What Is Claimed Is:

1. A method for dynamically allocating upstream bandwidth in a passive optical network that includes a central node and at least one remote node, wherein each remote node is coupled to at least one logical entity corresponding to a device or a user that transmits upstream data to the central node and receives downstream data from the central node, and wherein a shared out-going uplink couples the central node to an external network outside of the passive optical network, the method comprising:

receiving a request from a remote node for a grant to transmit upstream data from a logical entity associated with the remote node to the central node, wherein the size of the data to be transmitted does not exceed a transmission threshold assigned to that logical entity, and a logical entity may not request more than what is allowed by the corresponding transmission threshold;

if the request satisfies a bandwidth allocation policy, issuing a grant to the remote node to transmit upstream data;

in response to the grant, receiving upstream data from the remote node; and

transmitting the received upstream data to the out-going uplink according to a set of service level agreements.

2. The method of claim 1, wherein the transmission threshold assigned to a logical entity within a priority level is determined by considering:

the maximum allowable delay for that priority level;

data speed of the shared out-going uplink;

the logical entity's service level agreement; and

the total number of logical entities within that priority level.

3. The method of claim 1,

wherein prior to transmitting the received upstream data to the out-going uplink, the received upstream data is stored in a receiver buffer within the central node;

wherein the receiver buffer includes a number of FIFO queues, each of which buffers upstream data received from an associated logical entity; and

wherein transmitting the received upstream data involves retrieving and transmitting the upstream data stored in the receiver buffer to the out-going uplink according to a set of service level agreements.

4. The method of claim 3, wherein satisfying the bandwidth allocation policy requires that:

there is sufficient available space in the receiver buffer to accommodate the upstream data to be transmitted as requested; and

the logical entity from which upstream data transmission is requested is scheduled to transmit data next.

5. The method of claim 4, wherein all the logical entities within the passive optical network are scheduled to transmit upstream data using a hierarchical round-robin scheme by performing the following operations:

(1) grouping logical entities with the highest priority to form a top-priority level;

(2) allowing each logical entity in the top-priority level to transmit upstream data in a round-robin fashion by assigning a slot to each logical entity in the top-priority level;

(3) within the top-priority level, reserving one slot for lower-priority traffic;

(4) grouping logical entities with the second-highest priority to form a second-priority level;

(5) allowing each logical entity in the second-priority level to transmit data by assigning the reserved slot within the top-priority level to each logical entity in the second-priority level in a round-robin fashion;

(6) within the second-priority level, reserving one slot for lower-priority traffic; and

(7) repeating operations similar to operations (4) – (6) for logical entities with lower priorities until every logical entity is assigned a slot for transmitting upstream data according to its priority.

6. The method of claim 4, wherein all the logical entities within the passive optical network are scheduled to transmit upstream data using a strict priority scheduling scheme.

7. The method of claim 3, further comprising:

keeping a record of outstanding data for each logical entity, wherein outstanding data is upstream data for which a grant for transmission has been issued by the central node, but which has not been received by the central node;

subtracting the size of outstanding data from the unfilled space of the receiver buffer to calculate available space in the receiver buffer; and

after a period of time following issuance of a grant for transmitting a piece of data when the piece of data is due to arrive at the central node, removing the information for the piece of data from the record of outstanding data for to a logical entity, regardless of whether the piece of data has actually been received by the central node.

8. The method of claim 3, wherein retrieving and transmitting data stored in the receiver buffer to the out-going uplink according to a set of service level agreements involves retrieving and transmitting data stored in each FIFO queue using a hierarchical round-robin scheme in accordance with each logical entity's service level agreement.

9. The method of claim 1, wherein each remote node includes a number of queues, each of which is associated with a logical entity and stores upstream data from the device or user associated with that logical entity.

10. The method of claim 9,

wherein the request from a remote node reports the state of a queue within that remote node associated with a logical entity; and

wherein the request piggybacks on an upstream data transmission.

11. The method of claim 9, wherein if a FIFO queue within the receiver buffer in the central node is full, the issuance of grants to the corresponding logical entity is paused, thereby causing the queue associated with that logical entity within a remote node to become full, upon which the remote node generates a flow-control message to the corresponding device or user to slow down the upstream data transmission from that device or user.

12. The method of claim 1, wherein a remote node tracks the amount of time between the grants to transmit upstream data for each logical entity associated with the remote node; and

if the amount of time between grants exceeds a certain interval, the remote node sets an alarm and sends a message to the central node via an OAM frame,

whereby upon receiving the message, the central node is allowed to reset a record associated with the corresponding logical entity.

13. The method of claim 1,

wherein the central node periodically sends out polls to the remote nodes to see if a logical entity has any data to send;

wherein the polling frequency for a corresponding logical entity reflects the SLA of the logical entity; and

wherein if a non-poll grant has been previously sent to a logical entity, the subsequent poll to that logical entity is sent at a time after the non-poll grant, the time being calculated in accordance to the corresponding polling frequency.

14. The method of claim 13,

wherein a remote node tracks the amount of elapsed time between non-poll grants for each logical entity associated with the remote node;

wherein if the elapsed time between non-poll grants for a logical entity exceeds a certain interval, the remote node sets an alarm;

wherein if the alarm is set and the remote node has data to send from the corresponding logical entity, the remote node sends a message to the central node via an OAM frame denoting an error condition, which instructs the central node that the logical entity is in a locked up/hung state; and

wherein upon receiving the message, the central node is allowed to reset or modify a record associated with the logical entity.

15. An apparatus that dynamically allocates upstream bandwidth in a passive optical network, comprising:

a central node;

at least one remote node, wherein each remote node is coupled to at least one logical entity corresponding to a device or a user that transmits upstream data to the central node and receives downstream data from the central node;

a shared out-going uplink that couples the central node to an external network outside of the passive optical network;

a dynamic bandwidth allocation mechanism within the central node configured to,

receive a request from a remote node for a grant to transmit upstream data from a logical entity associated to the remote node to the central node, wherein the size of the data to be transmitted does not exceed a transmission threshold assigned to that logical entity, and a logical entity may not request more than what is allowed by the corresponding transmission threshold, and

if the request satisfies a bandwidth allocation policy, issue a grant to the remote node to transmit upstream data;

a receiving mechanism configured to receive upstream data from the remote node in response to the grant; and

a bandwidth shaping mechanism configured to transmit the received upstream data to the out-going uplink according to a set of service level agreements.

16. The apparatus of claim 15, wherein the transmission threshold assigned to a logical entity within a priority level is determined by considering:

the maximum allowable delay for that priority level;

data speed of the shared out-going uplink;

the logical entity's service level agreement; and

the total number of logical entities within that priority level.

17. The apparatus of claim 15, further comprising a receiver buffer within the central node configured to store the received upstream data;

wherein the receiver buffer includes a number of FIFO queues, each of which buffers upstream data received from an associated logical entity; and

wherein the bandwidth shaping mechanism is configured to retrieve and transmit the upstream data stored in the receiver buffer to the out-going uplink according to a set of service level agreements.

18. The apparatus of claim 17, wherein satisfying the bandwidth allocation policy requires that:

there is sufficient available space in the receiver buffer to accommodate the upstream data to be transmitted as requested; and

the logical entity from which upstream data transmission is requested is scheduled to transmit data next.

19. The apparatus of claim 18, wherein all the logical entities within the passive optical network are scheduled to transmit upstream data using a hierarchical round-robin scheme by performing the following operations:

(1) grouping logical entities with the highest priority to form a top-priority level;

(2) allowing each logical entity in the top-priority level to transmit upstream data in a round-robin fashion by assigning a slot to each logical entity in the top-priority level;

(3) within the top-priority level, reserving one slot for lower-priority traffic;

(4) grouping logical entities with the second-highest priority to form a second-priority level;

(5) allowing each logical entity in the second-priority level to transmit data by assigning the reserved slot within the top-priority level to each logical entity in the second-priority level in a round-robin fashion;

(6) within the second-priority level, reserving one slot for lower-priority traffic; and

(7) repeating operations similar to operations (4) – (6) for logical entities with lower priorities until every logical entity is assigned a slot for transmitting upstream data according to its priority.

20. The apparatus of claim 18, wherein all the logical entities within the passive optical network are scheduled to transmit upstream data using a strict priority scheduling scheme.

21. The apparatus of claim 17, further comprising an outstanding data tracking mechanism configured to:

keep a record of outstanding data for each logical entity, wherein outstanding data is upstream data for which a grant for transmission has been issued by the central node, but which has not been received by the central node;

subtract the size of outstanding data from the unfilled space of the receiver buffer to calculate available space in the receiver buffer; and

after a period of time following issuance of a grant for transmitting a piece of data when the piece of data is due to arrive at the central node, to remove the information for the piece of data from the record of outstanding data for to a logical entity, regardless of whether the piece of data has actually been received by the central node.

22. The apparatus of claim 17, wherein the bandwidth shaping mechanism is configured to retrieve and transmit data stored in each FIFO queue using a hierarchical round-robin scheme in accordance with each logical entity's service level agreement.

23. The apparatus of claim 15, wherein each remote node includes a number of queues, each of which is associated with a logical entity and stores upstream data from the device or user associated with that logical entity. 24.

The apparatus of claim 23,

wherein the request from a remote node reports the state of a queue within that remote node associated with a logical entity; and

wherein the request piggybacks on an upstream data transmission.

25. The apparatus of claim 23, wherein if a FIFO queue within the receiver buffer in the central node is full, the issuance of grants to the corresponding logical entity is paused; thereby causing the queue associated with that logical entity within a remote node to become full, upon which the remote node generates a flow-control message to the corresponding device or user to slow down the upstream data transmission from that device or user.

26. The apparatus of claim 15, wherein a remote node is configured to:

track the amount of time between the grants to transmit upstream data for each logical entity associated with the remote node; and

if the amount of time between grants exceeds a certain interval, to set an alarm and send a message to the central node via an OAM frame, whereby upon receiving the message, the central node is allowed to reset a record associated with the corresponding logical entity.

27. The apparatus of claim 15, wherein the central node is configured to periodically send out polls to the remote nodes to see if a logical entity has any data to send;

wherein the polling frequency for a corresponding logical entity reflects the SLA of the logical entity; and

wherein if a non-poll grant has been previously sent to a logical entity, the subsequent poll to that logical entity is sent at a time after the non-poll grant, the time being calculated in accordance to the corresponding polling frequency.

28. The apparatus of claim 27,

wherein a remote node is configured to:

track the amount of elapsed time between non-poll grants for each logical entity associated with the remote node;

if the elapsed time between non-poll grants for a logical entity exceeds a certain interval, set an alarm; and

if the alarm is set and the remote node has data to send from the corresponding logical entity, to send a message to the central node via an OAM frame denoting an error condition, which instructs the central node that the logical entity is in a locked up/hung state; and

wherein the central node is configured to reset or modify a record associated with the logical entity upon receiving the message.

29. A computer-readable storage medium storing instructions that when executed by a computer cause the computer to perform a method for dynamically allocating upstream bandwidth in a passive optical network that includes a central node and at least one remote node, wherein each remote node is coupled to at least one logical entity corresponding to a device or a user that

transmits upstream data to the central node and receives downstream data from the central node, and wherein a shared out-going uplink couples the central node to an external network outside of the passive optical network, the method comprising:

receiving a request from a remote node for a grant to transmit upstream data from a logical entity associated with the remote node to the central node, wherein the size of the data to be transmitted does not exceed a transmission threshold assigned to that logical entity, and a logical entity may not request more than what is allowed by the corresponding transmission threshold;

if the request satisfies a bandwidth allocation policy, issuing a grant to the remote node to transmit upstream data;

in response to the grant, receiving upstream data from the remote node; and

transmitting the received upstream data to the out-going uplink according to a set of service level agreements.

30. The computer-readable storage medium of claim 29, wherein the transmission threshold assigned to a logical entity within a priority level is determined by considering:

the maximum allowable delay for that priority level;

data speed of the shared out-going uplink;

the logical entity's service level agreement; and

the total number of logical entities within that priority level.

31. The computer-readable storage medium of claim 29,

wherein prior to transmitting the received upstream data to the out-going uplink, the received upstream data is stored in a receiver buffer within the central node;

wherein the receiver buffer includes a number of FIFO queues, each of which buffers upstream data received from an associated logical entity; and

wherein transmitting the received upstream data involves retrieving and transmitting the upstream data stored in the receiver buffer to the out-going uplink according to a set of service level agreements

32. The computer-readable storage medium of claim 31, wherein satisfying the bandwidth allocation policy requires that:

there is sufficient available space in the receiver buffer to accommodate the upstream data to be transmitted as requested; and

the logical entity from which upstream data transmission is requested is scheduled to transmit data next.

33. The computer-readable storage medium of claim 32, wherein all the logical entities within the passive optical network are scheduled to transmit upstream data using a hierarchical round-robin scheme by performing the following operations:

(1) grouping logical entities with the highest priority to form a top-priority level;

(2) allowing each logical entity in the top-priority level to transmit upstream data in a round-robin fashion by assigning a slot to each logical entity in the top-priority level;

(3) within the top-priority level, reserving one slot for lower-priority traffic;

(4) grouping logical entities with the second-highest priority to form a second-priority level;

(5) allowing each logical entity in the second-priority level to transmit data by assigning the reserved slot within the top-priority level to each logical entity in the second-priority level in a round-robin fashion;

(6) within the second-priority level, reserving one slot for lower-priority traffic; and

(7) repeating operations similar to operations (4) – (6) for logical entities with lower priorities until every logical entity is assigned a slot for transmitting upstream data according to its priority.

34. The computer-readable storage medium of claim 32, wherein all the logical entities within the passive optical network are scheduled to transmit upstream data using a strict priority scheduling scheme.

35. The computer-readable storage medium of claim 31, wherein the method for dynamically allocating upstream bandwidth in a passive optical network further comprises:

keeping a record of outstanding data for each logical entity, wherein outstanding data is upstream data for which a grant for transmission has been issued by the central node, but which has not been received by the central node;

subtracting the size of outstanding data from the unfilled space of the receiver buffer to calculate available space in the receiver buffer; and

after a period of time following issuance of a grant for transmitting a piece of data when the piece of data is due to arrive at the central node, removing the information for the piece of data from the record of outstanding data for to a

logical entity, regardless of whether the piece of data has actually been received by the central node.

36. The computer-readable storage medium of claim 31, wherein retrieving and transmitting data stored in the receiver buffer to the out-going uplink according to a set of service level agreements involves retrieving and transmitting data stored in each FIFO queue using a hierarchical round-robin scheme in accordance with each logical entity's service level agreement.

37. The computer-readable storage medium of claim 29, wherein each remote node includes a number of queues, each of which is associated with a logical entity and stores upstream data from the device or user associated with that logical entity.

38. The computer-readable storage medium of claim 37,

wherein the request from a remote node reports the state of a queue within that remote node associated with a logical entity; and

wherein the request piggybacks on an upstream data transmission.

39. The computer-readable storage medium of claim 37, wherein if a FIFO queue within the receiver buffer in the central node is full, the issuance of grants to the corresponding logical entity is paused, thereby causing the queue associated with that logical entity within a remote node to become full, upon which the remote node generates a flow-control message to the corresponding device or user to slow down the upstream data transmission from that device or user.

40. The computer-readable storage medium of claim 29, wherein a remote node tracks the amount of time between the grants to transmit upstream data for each logical entity associated with the remote node; and

if the amount of time between grants exceeds a certain interval, the remote node sets an alarm and sends a message to the central node via an OAM frame, whereby upon receiving the message, the central node is allowed to reset a record associated with the corresponding logical entity.

41. The computer-readable storage medium of claim 29,

wherein the central node periodically sends out polls to the remote nodes to see if a logical entity has any data to send;

wherein the polling frequency for a corresponding logical entity reflects the SLA of the logical entity; and

wherein if a non-poll grant has been previously sent to a logical entity, the subsequent poll to that logical entity is sent at a time after the non-poll grant, the time being calculated in accordance to the corresponding polling frequency.

42. The computer-readable storage medium of claim 41,

wherein a remote node tracks the amount of elapsed time between non-poll grants for each logical entity associated with the remote node;

wherein if the elapsed time between non-poll grants for a logical entity exceeds a certain interval, the remote node sets an alarm;

wherein if the alarm is set and the remote node has data to send from the corresponding logical entity, the remote node sends a message to the central node via an OAM frame denoting an error condition, which instructs the central node that the logical entity is in a locked up/hung state; and

wherein upon receiving the message, the central node is allowed to reset or modify a record associated with the logical entity.